

METHOD AND APPARATUS FOR BLOCKING EFFECT REDUCTION

This invention relates to the estimation of the subjective quality of a picture that has been decoded from a compressed bit-stream and in certain aspects to techniques for improving that subjective quality.

Generally, both the bitstream itself and the decoded picture will be accessible but the original source will not; hence the term 'single-ended' will be applied. Such a single-ended estimate will clearly not be as reliable as one in which the source picture can be compared to the decoded output, but it can serve as a useful indicator of potential problems in a broadcast chain involving compression when the bitstream is being monitored.

The invention relates in the most important example to the MPEG-2 video compression standard, but applies more broadly to transform based compression systems.

A problem to be solved is that of estimating the subjective picture quality of a picture or sequence decoded from an MPEG-2 bitstream. The usual method of performing such an estimate is referred to in this proposal as the "double-ended" method.

The decoded picture is compared with a necessarily delayed version of the source picture. The most common quality measure based on this comparison is the peak signal-to-noise ratio (PSNR) which is based on the ratio of the maximum possible signal power to the power of the difference between source and decoded signals. Other measures are available; some of which attempt to take into account human perception factors.

The disadvantage of all the double-ended methods is that they require access to the picture source. While this is appropriate for testing systems in a laboratory, it cannot normally be used for monitoring the quality of compression in the field.

The object of the present invention is to overcome that disadvantage by providing a series of quality estimation methods based on a "single-ended" approach.

- 2 -

Accordingly, the present invention consists in one aspect in a method of estimating the severity of a picture artefact arising from block based processing, comprising the steps of deriving a pixel difference signal and filtering the signal horizontally and vertically to derive a local measure of
5 artefact severity.

Advantageously, the measure is employed to control the operation of a filter adapted to conceal the visibility of said artefact.

Suitably, the measure is employed to control a fade between the picture signal and the output of said filter.

10 The present invention consists in another aspect in a method for estimating the signal to noise ratio of a picture signal decoded from a compressed bit-stream, comprising the steps of determining the quantization values employed in said compression and deriving said estimate by processing said values.

15 Preferably, a measure taken at an upstream location is passed forward for comparison with a measure taken at the device under test.

Embodiments of this invention make use of the "Information Bus" which is the subject of an earlier patent applications (see EP 0 765 576 and EP 0 913 058) . The Information Bus is a signal containing all the compression
20 coding decisions and parameters extracted from the compressed bitstream, in an easily accessible form. More sophisticated versions of the quality estimation techniques presented here may also make use of the information signal which is also the subject of EP 0 807 356. This is similar to the Information Bus but carries information about other processes that may have
25 taken place upstream of the compression codec under consideration.

This invention will now be described by way of example with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a first embodiment of the present
30 invention;

- 3 -

Figures 2 and 3 are block diagrams of picture quality measures for use in the arrangement of Figure 1;

5 Figure 4 is a block diagram of a further embodiment of the present invention; and

Figure 5 is a diagram of a filter for use in the arrangement of Figure 4.

10 The basic architecture of single-ended quality measurement is shown in Figure 1.

15 The MPEG signal from a remote, upstream encoder 100 is received by a decoder 102 which is adapted as described in the previously mentioned references to provide an Information Bus output, in addition to the decoded video. The picture quality measurement unit 104 therefore has access not only to the decoded video but also to coding decisions which were taken by the encoder 100 and which are of course implicit in the MPEG bit-stream.

20 The picture quality measurement process in one embodiment operates only from information available at the decoder side of the compression codec; the decoded video signal and the Information Bus containing the coding decisions and parameters. It has no access to the picture source. Because of this, the quality measurement can never be completely reliable because there is no way of telling which degradations in the picture are due to the current coding process and which were on the source. So it is not intended as a full replacement for laboratory measurements based on the double-ended approach. But it is useful for many applications and it is certainly
25 sufficient for those monitoring applications in the field where a simple automatic indication of the "red - amber - green" variety is required. However, there will be described later a further embodiment of the invention which involves a modification by which some account can be taken of the source
30 material.

One of the most frequent complaints about MPEG-2 coded pictures is that they appear "blocky", meaning that the block and macroblock structure of the picture is visible. These blocking artefacts can occur for several reasons:

- 5 • Variation in quantizer scale between macroblocks
- Coarse quantization of DC coefficients in non-intra macroblocks
- Residual visibility of a prediction error resulting from a non-uniform motion vector field

10 Instead of attempting to analyse each of those possible causes, the "blockiness" measure provided in this embodiment is based simply on the end result, i.e. the decoded picture. There are various possible measures of blockiness, but the principle behind all of them is to compare pixel differences across block boundaries with pixel differences not across block boundaries.

15 In the discussion that follows, care should be taken to recognise the distinction between *macroblock* (16x16 block) boundaries and *DCT block* (8x8 block) boundaries.

The following is an example of a measure of blockiness that works on macroblock boundaries:

20 *Horizontal macroblockiness* = the picture-by-picture mean absolute horizontal adjacent luminance pixel difference across macroblock boundaries, expressed as a fractional increase over the mean absolute horizontal adjacent pixel difference not across DCT block boundaries

25

An example showing how this measure could be implemented in hardware is given in Figure 2.

Pixel differences are taken across a pixel delay and the absolute value calculated. The result is fed to two gated accumulators controlled by a modulo-16 pixel counter which is reset by a line synchronization pulse. The

30

- 5 -

upper accumulator sums the pixel differences across macroblock boundaries (when the modulo-16 pixel count = 0) and the lower accumulator sums the pixel differences not across DCT block boundaries (when the modulo-16 pixel count \neq 0 or 8). Event counters count the occurrences of each of these two cases so that the dividers can calculate mean values of the two quantities. Finally, the fractional increase is calculated, giving the blockiness measure. The accumulators and event counters are reset once per picture.

This particular measure has the interesting property that, when applied to frames that were I-frames in the MPEG-2 bitstream, the result is almost exactly proportional to the average quantizer scale value. When applied to P and B-frames, the result is smaller but reflects quite clearly differences in perceived blockiness arising from differences in motion estimation systems.

The following variations in the definition of the blockiness measure are possible and are considered to be part of the invention:

- The DCT block boundary can be used instead of the macroblock boundary. This would require a change the logical outputs of the pixel counter. Note that in both this and the original case the denominator of the fraction is the pixel difference not across DCT block boundaries.

- The difference could be taken vertically rather than horizontally (requiring a line delay instead of a pixel delay), or a combination of the two could be used. We have chosen the horizontal difference because this is much easier to calculate in hardware and because the boundaries are the same whether field or frame picture coding was used. However, they may be circumstances in which the vertical differences are easier to calculate.

- Mean square values, or the mean of some other function of pixel differences, could be used instead of mean absolute differences.

- 6 -

5 . Some statistical function other than the mean could be used. For example, because it might be considered that very poor blockiness in a small region of the picture might be more disturbing to the eye than an evenly distributed blockiness resulting in the same average value, it might be better to use, for example, the 90th centile of the macroblock boundary pixel difference.

10 . The blockiness could be expressed as a logarithmic ratio (like a dB measure) rather than a fractional increase. This would affect the final block in Figure 2. It is noted however that a straight difference has been shown experimentally to give good results.

15 . It may be possible to use a reduced number of pixel differences in the measure.

20 . The measurement period could be greater or less than one picture period. This would affect the resetting of the accumulators and event counters in Figure 2.

25 It is desirable to record the blockiness separately for I-frames, P-frames and B-frames. The figures are much lower in P and B-frames because the denominator of the expression contains prediction residues that may have come from macroblock or block boundaries in reference frames. To detect the picture type (I, P or B), the Information Bus could be used. Alternatively, in the absence of the Information Bus, a method of picture type detection such as that described in the contemporaneously filed patent application (claiming priority from GB 9822092.4) could be used. A further possibility is that the variations in the blockiness measure itself could be used as the basis of a method of picture type detection.

- 7 -

The above description assumes that the positions of the macroblock boundaries are known. In some cases, this information may not be available. However, it is possible to obtain this information by calculating the blockiness assuming each of the 16 possible positions in turn (either in full or using a reduced number of pixels) and choosing the position that yields the maximum value.

It will be useful in certain applications to have a measure of blockiness, not just for a picture, but also for regions of the picture.

An alternative embodiment for deriving a measure of blockiness will now be described, with reference to Figure 3.

Delay 300 and subtractor 302 serve to generate the pixel difference signal. The absolute value is taken at 304 and the result is filtered at 306. In the case of 8 x 8 blocks, the aperture for this filter might be:

-1 -1 -1 14 -1 -1 -1

so as to produce directly the mean absolute difference across macroblock boundaries relative to the mean not across macroblock boundaries. The result is then sub-sampled by 8 in block 308, to remove the meaningless values where the filter is not centred on a macroblock boundary.

The functions of blocks 310 and 312 are, similarly, to construct a vertical average over the 8 lines of a macroblock and then to discard the values which are not required. The resulting measure can be averaged over an entire picture as at 314 or filtered in local filter 316 (with typically an 8 x 8 moving average) to provide a local measure of blockiness.

It is possible to use the local and global measures in an ingenious fashion to provide for a filter to remove blockiness, which varies in effect in accordance with the local blockiness.

Reference is directed to Figure 4 which shows a median filter 402, selecting the middle-ranking of three inputs. One input is the uncorrected video input from which has been subtracted at 402 a combined blockiness measure from combiner 404. This receives both the global and local

- 8 -

blockiness measures and produces a combined measure. The second median filter input is the sum of the combined blockiness measure and the uncorrected signal whilst the third input is the output of a filter constructed to remove those horizontal and vertical spatial frequencies associated with blockiness. This filter might have pass bands as shown in Figure 5.

The operation of the described arrangement is to provide an arithmetic fade between the uncorrected video and the filter output, in dependence on the combined blockiness measure. The blockiness measure is effectively compared with the absolute difference between the uncorrected video and the filter output. If the blockiness measure exceeds this difference, the full filtered output is employed. Where the blockiness measure lies between zero and this difference, an arithmetic fade is provided.

A further single-ended quality measure provides an estimate of the peak-signal-to-noise-ratio (PSNR) of the decoded picture. It has been found surprisingly that by using essentially only the quantizer scale values present in the bit-stream an estimate of PSNR can be produced which corresponds well with actual experimental values.

The approach taken is as follows:

$$PSNR_{estimate} = PSNR_0 - \frac{\sum_{allmacroblocks} \log q}{N_{macroblocks}} (A + BH)$$

20

where q is the the quantizer level spacing q , which is known from the quantizer scale code and q_scale_type parameters received in the Information Bus and A and B are experimentally derived parameters. The base ratio $PSNR_0$ is taken as an experimental value of signal to noise ratio employing the finest allowable quantisation and a pre-determined quantisation weighting matrix

25

The sum provides an average quantiser scale value for the picture, the average being taken in the logarithmic domain over all the macroblocks in the picture.

- 9 -

H is an estimate of the entropy of the most recent I-picture which is defined here as the number of DCT coefficient bits that would be generated with a quantiser_scale value of 1, the finest possible. The estimate of H is based on the actual number of coefficient bits generated for each macroblock and the quantiser_scale value used for the macroblock, using a quadratic formula based on a model.

Thus

$$H = \sum_{\text{all macroblocks}} cf(Mq)$$

where c is the number of coefficient bits in the macroblock. The function of Mq is preferably quadratic.

M is a correction factor which corrects for the use of different weighting matrices in the bitstream. It is defined such that it has the value 1 when the matrices used are the MPEG default intra matrix and the MPFG default (i.e. flat) non-intra matrix.

Finally, a correction involving the number of displayed and encoded samples may be added in the form

$$+ 10 \log \frac{N_{\text{coded}}}{N_{\text{displayed}}}$$

This is designed to correct for the subjective effects of working with subsampled pictures and also with letterbox pictures. For example, if a 720 x 576 picture containing a 720 x 332 letterbox is sub-sampled to 480 x 288 for encoding (so that only 166 of the coded lines contain active picture data), the value of N_{coded} is 480 x 166 and the value of $N_{\text{displayed}}$ is 720 x 576. Note that this correction is for subjective effects and should not be applied if a correlation with PSNR calculated over the whole of the coded picture is being sought.

The approaches described thus far are based on the "single-ended" architecture and as such suffer from the limitation that there is no knowledge of how much of the impairment being measured has come from the coding

- 10 -

process and how much has come from the source. A modification will now be described in which that limitation can be partially overcome.

The technique is to apply some or all of the measures to the source and/or to intermediate points in the signal chain and to transmit the results to the decoder under consideration, using a combination of ancillary data in MPEG bitstreams and the Information Bus, according to the principles of the EP 0 807 356. At intermediate points in the chain, where the picture has been decoded from an MPEG bitstream and there is access to the Information Bus resulting from that decoding process, all the measures described above can be used. At the source, or at places where a full Information Bus is not available, the choice of measures may be more limited. In either case, the results can be compared with the current results and the difference will give an indication of how much of the finally measured degradation was due to the intervening compression process or processes.

It should be understood that this invention has been described by way of example only and that a wide variety of modifications are possible without departing from the scope of the invention.